

Welcome to the REX Ops Group

Adam Lyon

Fermilab / Computing Division

December 7, 2006

Notes from meeting in
these yellow boxes

WELCOME!!!

REX Ops Group

- Adam (Group Leader)
- Angela (joins 2007.01.02)
- Chuck
- Dehong
- Krzysztof
- Randolph
- Robert (Assistant Group Leader)
- Steve
- Perhaps one or two CD postdocs (future)

Up the chain

- Qizhong Li (REX Department Head Deputy)
- Rick Snider (REX Department Head)
- Patty McBride (Scientific Computing Head)
- Vicky White (Division Head)

Not clear if we are really getting postdocs

Basic Message for today

Keep doing your regular duties, unless I've told you otherwise.

- December is a bad month to make changes
 - For me, DØ collaboration meeting
 - I will be away for most of next week
 - Holidays
- We'll start implementing changes mentioned here in January.

Responsibilities

Our responsibilities are to maintain smooth running for:

- Data Handling at CDF (includes dCache, SAM, and remote services)
- Data Handling at DØ (includes SAM and SAMGrid)
- Data Handling at MINOS (SAM and dCache)
- Online tape logging for CDF
- Online data logging and tape logging for DØ
- Software builds for CDF
- Plan for future needs with stakeholders (experiments)
- Special projects and improvements (MIS, SQL Builder)

DBB is a part of "data handling"

What kinds of things do we do?

- Monitor systems (we try to find the problems first)
- Field problems and questions from shifters and users
- Troubleshoot to determine responsible party
- Coordinate fix/work-around with responsible party
- Coordinate long term fixes (e.g. down times)
- Do light development for fixes and improvements
- Coordinate and plan with experiments
- Do some longer term development to improve operations

Who is responsible for what?

There are several options. Let's go through a few...

Each is responsible for only a particular system all the time

PROS:

- You get to become a true expert on "your" system
- Very efficient use of manpower
- Customer support administration is easy (you handle it for your system)

CONS:

- Vacation and leave are very problematic
- Group becomes highly compartmentalized
- Perhaps quick burn out
- You must watch customer support mailing lists all the time

The SAM team used to be like this structure.

Each is responsible for everything all the time

You know everything about all the systems.

PROS:

- Each system has lots coverage
- Anyone can field questions/problems
- Lots of experts

BUT: Some overlap of knowledge is extremely desirable

CONS:

- Too much to know
- Huge duplication of manpower
- Administratively difficult
- Who is responsible for watching the lists? [One person? But that's a big job to have all the time.]

Compromise: Focus knowledge but not exclusive

Certain people have main knowledge for a particular system, but is also shared among the group. Also rotate primary support role.

PROS:

- Concentrate on system
- Crucial knowledge is shared
- Vacations/Leave are not problematic
- We all share the primary support load, but not everyone all of the time

CONS:

- Need documentation
- Need organization
- Potential heavy load for primary support person
- More practical, but still lots to know
- Unless good docs, still need experts

Rick comments that the "compromise" scheme can become like extreme scheme 1 if we don't really spread knowledge. Documentation is the key and all will be expected to contribute.

What does this mean?

- You will focus on a particular system or aspect (perhaps more than one and perhaps with others)
- You may do some light development on the system (e.g. improve monitoring, operational load, fixing scripts)
- It is in your best interest to write documentation so that others can troubleshoot/cleanup/fix without bothering you
- Primary support role will rotate to you
- Secondary support role will rotate to you
- **Important:** While your focus may be on a system for a particular experiment, your support role is experiment neutral

Perhaps you are secondary first to ramp up into primary the next week

Support Roles

Primary:

- You must watch lists, watch monitoring, answer cell phone
- You are responsible for fielding problems and questions (at a priority above your normal duties)
- You are responsible for resolving the question/problem or finding the person who can
- You must stay on the issue until it is resolved (by you or someone else) or until your shift rotates away
- You assume ownership of issues not resolved by the previous shift
- Because this can be a high load week, the secondary assumes your role when you are unavailable (abducted by aliens) or are overwhelmed. You should coordinate with secondary

Secondary:

- Provide help and backup to the primary

Tools we'll need

Documentation SAM uses lots of free standing web pages and Plone (with a Wiki). Plone is slow and not easy, but is supported

Issue Tracker SAM uses the Plone Issue Tracker. Not plugged into helpdesk. Supported at the hobby level. May have to go to Remedy

Task Tracker Jason Allen wrote a task tracker to track long term tasks for RunII-sys. Perhaps we use that too

Code Repository CDCVS?

No solution is perfect. Goal is to choose something we do not have to support.

Is plone secure? - yes, can write "private" pages that only our group can see

Things you should think about

Our success depends on simplicity and efficiency

- Can I make this system easier to operate?
 - Replace manual procedures with scripts
 - Automate checks (e.g. cab checker)
 - Improve the system
- Can I make this system easier to monitor?
 - Consolidated monitoring (e.g. Dehong's)
 - Automated alerts (hard to do right) NGOP?
- Can I make this system easier to recover?
 - Hardware failures happen (often)
 - Improve/automate recovery (e.g. nameserver)

Perils of an Operations Group

When things are going right, nobody notices us.

- Have metrics and reports where we can ...
 - 1 Quantify how things are going right
 - 2 Quantify what we've done to get there
 - 3 Publicize our successes

When things go wrong, everybody notices us.

- Eventually, it's going to happen
- DØ and CDF represent two extremes
- If major unplanned downtime, we instigate root cause analysis with experiment
 - SAM never really did this; but should have
 - Don't hide our mistakes; learn from them

Randolph points out that "root cause" is too restrictive. Should also include ancillary problems that were exposed by the "root" problem.

Meetings We will have an open operations meeting every week (like we do now that Krzysztof has been running). Perhaps also have a closed group meeting as necessary

Time Sheets Give time sheets to me or Robert to sign (perhaps at the open Ops meeting?)

Reports Send weekly reports to myself and Robert. Since we're an operations group, weekly reports is the best fit

Vacations/Leave Use the CD leave request system to plan vacation. If you take a sick or unplanned vacation day . . .

- 1 Contact myself or Robert ASAP (e-mail or phone)
- 2 Request that leave in the CD system (that day if possible, or the day you return to work)

Goals Expect to have meetings about goals late this month and early January. I will try to do quarterly review of goals

Conclusions

- While small, we are a highly expert group. Each of us brings an excellent record of experience and reliability
- The tasks ahead are difficult, but I'm confident that **working together** we can succeed
- Don't keep ideas, problems, achievements, complaints to yourself
- My job is to enable all of you to be successful; never hesitate to come to me if you need to